

HP OpenView

Storage Mirroring application notes

High availability for Microsoft SQL Server 2000

Legal and notice information

© Copyright 2005–2006 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Storage Mirroring High availability for Microsoft SQL Server 2000 application notes

Document overview

This document is a Storage Mirroring application note. An application note provides guidelines on the use of Storage Mirroring in a specific environment.

This document contains:

- **Document overview**—Explains what an application note contains, how it should be used, what you need to know before trying to use the application note, and where you can go for more information.
- **Solution overview**—Explains how the application works with Storage Mirroring and describes the considerations that you must weigh when implementing your Storage Mirroring solution. Review this section to make sure that you understand the theory involved with using Storage Mirroring and your application. Includes both basics, such as system requirements, as well as configuration and environment-specific topics, such as interactions with specific clients or special considerations for WAN (Wide Area Network) environments. Pay special attention to those topics that are directly related to your environment.
- **Sample implementations**—Describes specific examples of how to use Storage Mirroring for this solution. This includes information about the specific system setup used in the sample implementation. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Audience

This document is written for network and application administrators who have a working understanding of the applications and environments where the Storage Mirroring solution is to be deployed. You may need to expand on the documented information in order to customize the solution to fit your environment.

Before you use this application note, you should have an understanding of:

- Storage Mirroring
- SQL Server

Expectations

Application notes are intended to provide a framework for configuring a Storage Mirroring solution in a specific environment and to draw attention to decisions you will need to make when configuring your solution.

Because there are an infinite number of possible configuration, network, and environment scenarios, application notes contain general configuration guidelines as well as an example configuration procedure that has been tested for a specific environment.

This document assumes that you are comfortable working with your operating system, Storage Mirroring, and SQL Server.

Related documentation

Before you begin to configure your solution, make sure that you have complete documentation for your operating system, application, and Storage Mirroring. This application note does not provide step-by-step instructions for using standard operating system, application, and Storage Mirroring functionality.

The following documents contain additional information that you may need while setting up this solution:

- *HP OpenView Storage Mirroring user guide* or online documentation
- Reference guides or documentation for SQL Server

Getting help

Hewlett-Packard has application notes that describe how to configure Storage Mirroring with a variety of popular third-party applications. These application notes are available on the Storage Mirroring web site: <http://h18006.www1.hp.com/products/storage/software/sm/index.html>.


For help using Storage Mirroring, refer to the Storage Mirroring online manual or online help.

Solution overview

Microsoft SQL Server is a scalable, reliable, flexible, and high-performance relational database management system for Microsoft Windows 200x server-based systems. Storage Mirroring provides real-time enterprise data protection and replication. Storage Mirroring can be used to provide high availability for your SQL Server.

This document describes the steps necessary to configure Storage Mirroring to provide high availability for Windows 200x servers running Microsoft SQL Server 2000. These procedures allow a secondary server to assume the identity and role of a failed SQL Server while maintaining the availability of SQL services with minimal disruption or data loss.

To complete these instructions, you will install SQL Server 2000 and Storage Mirroring, and configure Storage Mirroring for replication and failover. Due to the complexities of these applications, this document is intended for network administrators with experience installing, configuring, and maintaining network applications including Storage Mirroring and Microsoft SQL Server.

 **NOTE:** Storage Mirroring allows you to configure one target to monitor and failover for one or more source machines. In a one-to-one configuration, you will want to replicate your SQL data to the same location on the target so that failover is automatic. In a many-to-one configuration, each SQL data store will need to be replicated to a unique location and then renamed to the corresponding SQL directory on the source before failover occurs.

How to use this application note

Because it is common for a SQL environment to include several front-end clients that attach to multiple databases (for example, one for entering customer information, one for processing incoming inventory, one for processing outgoing inventory, and so on), this application note includes instructions for how to configure high availability for SQL for multiple scenarios, including:

- High availability and disaster recovery between SQL 2000 standalone servers, described in "[One-to-one failover for SQL 2000 servers](#)" on page 5
- High availability and disaster recovery from a cluster to a standalone server, described in "[Cluster-to-standalone failover](#)" on page 12
- High availability and disaster recovery from a cluster to another cluster, described in "[Cluster-to-cluster failover](#)" on page 21

Refer to the appropriate section for information about configuring SQL 2000 failover for your environment.

Modifying the sample script files

After you modify the sample scripts, save them with a new name to remove the `SAMPLE_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

Configuring the replication set

For a default SQL installation, you *must* include the SQL Data directory in your replication set (`<drive>:\Program Files\Microsoft SQL Server\MSSQL\Data`). This directory includes the SQL application data and transaction log files.

In addition, you *may* need to include the following directories, depending on your environment configuration:

- Tempdb files. You should always include the `tempdb` files, unless you can determine that they are not being used by any application. Some applications, including Prophecy, PeopleSoft, and BizTalk®, write data to the `tempdb` file.
- Any other directories (even if on different drives) that you may have created to store SQL data files.

- The SQL Log directory (<drive>:\Program Files\Microsoft SQL Server\MSSQL\Log). This directory includes the SQL error logs from the source, which may be useful when determining the cause of a source failure. In order to prevent overwriting the target's error logs, you will need to replicate the log files to a different location on the target. Create a separate replication set for the error logs and select the "All to One" mapping option.

You *do not* need to replicate the application files, since they already exist on the target machine.

Enabling compression

By enabling compression, you can reduce the amount of bandwidth needed to transmit Storage Mirroring data. When compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it uncompresses it and then writes it to disk. On a default Storage Mirroring connection, compression is disabled.

Because the files that should be included in a SQL replication set can generate a significant amount of data, you should enable compression for the connection. For more information about enabling compression, see the *HP OpenView Storage Mirroring user guide*. However, keep in mind that the process of compressing data impacts processor usage. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled.

Configuring SQL memory usage

Storage Mirroring uses memory to queue operations and data on both the source and target. Since the source server is typically running a production application, it is important that the amount of memory Storage Mirroring and the other applications use does not exceed the amount of RAM in the system. If the applications require more memory than there is RAM, the system will begin to swap pages of memory to disk and the system performance will degrade.

For instance, SQL Server will use all of the available system memory when needed by default, and it may use almost all of the system memory during high-load operations. These high-load operations are precisely what cause Storage Mirroring to need memory to queue the data being changed by SQL Server. On a server with 1GB of RAM running SQL Server and Storage Mirroring, you might configure SQL Server to use only 512MB and Storage Mirroring to use 256MB, leaving 256MB for the operating system and other applications on the system. Many other server applications will use almost all system memory by default, so it is important to check and configure applications appropriately, particularly on high-capacity servers.

One-to-one failover for SQL 2000 servers

This section describes an example of how to configure Storage Mirroring and SQL Server. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Storage Mirroring allows you to configure one target to monitor and failover for one or more source machines. In a one-to-one configuration, you will want to replicate your SQL data to the same location on the target so that failover is automatic.

This section focuses on a single SQL Server being replicated to a single target.

In a many-to-one configuration, each SQL data store will need to be replicated to a unique location and then renamed to the corresponding SQL directory on the source before failover occurs.

Requirements

For this scenario, your system must meet the following requirements:

- Microsoft Windows 200x with the latest service pack
- Licensed copies of Microsoft SQL Server 2000 with the latest service pack
- Licensed copies of Storage Mirroring with the latest service pack

Install software on the source

1. Install SQL Server 2000 on the source, if it is not already installed.
2. Record the drive and directory where SQL is installed. The default directory for SQL 2000 is
<drive>:\Program Files\Microsoft SQL Server\MSSQL

SQL Installation Drive and Directory: _____

3. Install Storage Mirroring on the source machine using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.

Install and configure software on the target server

1. Install Storage Mirroring on the target using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.
2. Select **Start, Programs, Administrative Tools, Services** and double-click the Storage Mirroring service.
3. Select the Log On tab. Select the **Allow Service to Interact with Desktop** check box and click **OK**.
4. Install SQL 2000 on the target using the same drive and directory specifications recorded in step 2 of the previous section.
5. Set the following services to manual startup so that all SQL files are closed on the target and the Storage Mirroring source can replicate the changes. Make sure that you set all of the services, SQL included, relevant to your environment.
 - Distributed Transaction Coordinator
 - Message Queuing
 - MSSQLServer
 - SQLServerAgent
 - Microsoft Search




NOTE: If a failure should occur, the failover and failback scripts that you will be creating will control the stopping and starting of the SQL services on the source and target.

Configure and begin mirroring and replication

1. Select **Start, Programs, Storage Mirroring, Management Console**.
2. Double-click your source machine to log on.
3. Right-click on the source and select **Properties**.
4. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
5. Right-click your source machine and select **New, Replication Set** and enter the desired name for the replication set.
6. Select the SQL data you wish to protect. For more information about what you should include in your SQL replication set, see "[Configuring the replication set](#)" on page 4
7. Right-click the replication set name and select **Save** to save the replication set.
8. Drag and drop the replication set onto the target. The Connection Manager dialog box opens.
9. The **Source Server**, **Target Server**, **Replication Set**, and **Route** fields will automatically be populated. If you have multiple IP addresses on your target, verify the **Route** field is set to the correct network path. (For detailed information on connecting a source and target, see the *HP OpenView Storage Mirroring user's guide*.)
10. Select **One to One** to map the replication set data from the source to an identical volume/directory structure on the target.
11. On the Orphans tab, select the **Move/Delete Orphan Files** checkbox.
12. On the Mirroring tab, select the type of mirror, either **Full** or **File Differences**, to perform.

NOTE: If the target has been previously mirrored to or restored, select **File differences** with the **Use block checksum** option so that only the changed data is sent across the network.


13. Click **Connect** to start the mirror and replication processes.

 **NOTE:** If you start SQL Server and mount the replicated databases on the target, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. If the updated data on the target is needed, restore the data from the target to the source. Configure failover and begin failure monitoring

The following steps should be performed on the target machine.

1. Use the HP NSISPN utility to modify the SQL Service Principal Names (SPNs) of the source and target servers when failover and failback occur. The Storage Mirroring service account must also be given the appropriate permissions to modify the SPNs. Create a batch file called `prefailover.bat` using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

If you **are** using the same domain user account on the source and target servers, you do not need to update SPN permissions. Proceed to the next step

 **NOTE:** After you modify the sample scripts, save them with a new name to remove the `sample_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

sample_prefailover.bat

```
rem ***SAMPLE*** pre-failover script for one-to-one failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem When the SQL Server service is configured to use different domain user accounts
rem on the source and target servers, use commands similar to the following.
rem The default SQL Server port, 1433, is used in these examples.
rem SourceSqlHost refers to the SQL server name on the source server.

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlServiceAccount
NSISPN -D MSSQLSvc/SourceSqlHost:1433 SourceSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost:1433 TargetSqlServiceAccount

rem When the SQL Server service is started with the server's System account
rem use commands similar to the following.

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlHost
NSISPN -D MSSQLSvc/SourceSqlHost:1433 SourceSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost:1433 TargetSqlHost
```

2. If a failure occurs, you will want the SQL services to start on the target machine automatically. To do this, create a batch file called `postover.bat` using the sample batch file below. The following script should be run after failover has completed.

sample_postover.bat

```
rem ***SAMPLE*** post-failover script for one-to-one failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem You may need to remark out some commands depending on the function of your target.
rem The following commands should be used if your target is functioning as a back-end
rem Make sure that you start all of the services, SQL included, relevant to your
rem environment.

net start "Distributed Transaction Coordinator"
net start "MSSQLSERVER"
net start "SQLSERVERAGENT"
net start "Microsoft Search"
```

3. After a failure is resolved, you will be ready to bring your source back online. At this time, you will want to stop the SQL services on the target automatically. To do this, create a batch file called `preback.bat` using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

If you are using the same domain user account on the source and target servers, you do not need to update SPN permissions. Remove the NSISPN commands in the sample script.

sample_preback.bat

```
rem ***SAMPLE*** pre-failback script for one-to-one failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem You may need to remark out some commands depending on the function of your target.
rem The following commands should be used if your target is functioning as a back-end.
rem Make sure that you start all of the services, SQL included, relevant to your
rem environment.

net stop "Distributed Transaction Coordinator"
net stop "MSSQLSERVER" /y
net stop "SQLSERVERAGENT"
net stop "Microsoft Search"

rem When the SQL Server service is configured to use different domain user accounts
rem on the source and target servers, use commands similar to the following.
rem The default SQL Server port, 1433, is used in these examples
rem SourceSqlHost refers to the SQL server name on the source server.

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlServiceAccount
NSISPN -D MSSQLSvc/SourceSqlHost:1433 TargetSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost:1433 SourceSqlServiceAccount

rem When the SQL Server service is started with the server's System account
rem use commands similar to the following in the failback scripts:

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlHost
NSISPN -D MSSQLSvc/SourceSqlHost:1433 TargetSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost:1433 SourceSqlHost
```

Set SPN permissions

Service Principal Name (SPN) is a property of Active Directory computer and user accounts, and the Storage Mirroring service must have the appropriate permissions to modify SPNs on the source's SQL Server service account. Accordingly, the `WriteServicePrincipalName` permission on the SQL Server's service account in Active Directory must be assigned to the Storage Mirroring service account.

The following procedure includes references to the SQL Server service account and the Storage Mirroring service account. Use the following guidelines to determine which specific Active Directory user or computer account to use:

- If a service is using a user account, then that user account should of course be used when the service account is referenced.
- If a service is using the Local System account, then the computer account should be used when the service account is referenced.

To find what account a service is using, run the Windows Services snap-in and check the Log On tab on the service's Properties dialog to see if the service is using a user account or the Local System account.

Use the following procedure to assign the SQL Server service's `WriteServicePrincipalName` permission to the Storage Mirroring service account:

1. Click **Start, Run**, and type **MMC**. Click **OK**.
2. Select **File, Add/Remove Snap-in**.
3. Click **Add** and select **ADSI Edit**.
4. Click **Add**. Click **Close**, then click **OK**.
5. Right-click **ADSI Edit** and select **Connect To**.
6. Type in your domain name and click **OK**.
7. Expand the tree in the left pane to the `CN=Users` folder.
8. In the Users folder, locate the source's SQL service account.
9. Select the target's Storage Mirroring service account and click **View/Edit**.
10. Select the Security tab, click the **Advanced** button.
11. Click **Add**, then click **Object Types**.
12. Select any unchecked check boxes and click **OK**.
13. Type in or locate the target server's Storage Mirroring service account and click **OK**.
14. Select the Properties tab and select **WriteServicePrincipalName**.
15. Click **OK** to accept the change, then click **Close**.

Failure monitoring

1. On the target, select **Start, Programs, Storage Mirroring, Failover Control Center**.
2. Select the target machine from the list of available machines. If the target you need is not displayed, click **Add Target**, enter the machine name, click **OK**, and then click **Login**.
3. To add a monitor for the selected target, click **Add Monitor**. Type the name of the source machine and click **OK**. The Monitor Settings window will open.
4. In the left pane of the Monitor Settings window, mark the IP address that is going to failover. Verify that the correct NIC is selected in the right-hand pane under Current IP Address(es).
5. Highlight the machine name under Names to Monitor and click **Scripts**. Specify the scripts that were created earlier using `prefailover.bat` for the target pre-failover script, `postover.bat` for the target post-failover script, and `preback.bat` for the target pre-failback script.
6. Click **OK** to go back to the Monitor Settings dialog box and then click **OK** to begin monitoring the source machine.

In the event of a source machine failure, your target machine is now ready to stand in for your source. For information on monitoring failover, see the *HP OpenView Storage Mirroring user's guide*.

Dealing with a failure

In the event that the source server becomes unavailable, the following will occur:

1. The Failover Control Center will count down to zero, and a dialog box will appear on the target server indicating the failure. You may initiate failover by clicking **OK**.
2. After clicking **OK**, the target server will assume the IP address and server name of the source server.
3. The failover scripts will execute, making the necessary changes to the SPNs and starting the SQL services on the target server.
4. Use `NSISPN -L SourceSqlHost` and `NSISPN -L TargetSqlHost` to verify that the changes were made.

 **NOTE:** The Windows Resource Kit SETSPN utility can also be used in place of the NSISPN utility.

The target server is now standing in for the source SQL server. Clients may access SQL services on the target server just as if it were the original source.

Restoring your SQL data

If your source experiences a failure, such as a power, network, or disk failure, your target machine will stand in for the source while you resolve the source machine issues. During the source machine downtime, data is updated on the target machine. When your source machine is ready to come back online, the data is no longer current and must be updated with the new data on the target machine.

1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.


 **NOTE:** If you must rebuild your hard drive, continue with step 3.

If you do not need to rebuild your hard drive, verify that the Storage Mirroring connection on the source has been disconnected (right-click the connection in the Storage Mirroring Management Console and select **Disconnect**) and then continue with step 6.

-
3. Install Windows. Since your source machine is not connected to the network, go ahead and use the source's original name and IP address.
 4. Install Storage Mirroring using the installation defaults.
 5. Install SQL using the same drive and directory settings recorded in step 2 of "[Install software on the source](#)" on page 6.
 6. On the source, **stop** the following services (and any other SQL services relevant to your environment). Verify that all SQL files are closed on the source so that the Storage Mirroring target can replicate the changes back to the source. Depending on the type of failure, your services may be set to manual startup but could still be running
 - Distributed Transaction Coordinator
 - Message Queuing
 - MSSQLServer
 - SQLServerAgent
 - Microsoft Search
 7. On the target, select **Start, Programs, Storage Mirroring, Failover Control Center**.
 8. Select the target machine that is currently standing in for the failed source.
 9. Select the failed source and click **Failback**. The pre-failback script entered during the failover configuration stops the SQL services on the target so that no additional changes can be made.
 10. You will be prompted to determine if you want to continue monitoring the source server. Do not choose **Continue** or **Stop** at this time.
 11. Connect the source machine to the network.
 12. After the source is back online, select whether or not you want to continue monitoring this source machine (**Continue** or **Stop**).

 **NOTE:** Verify that the Storage Mirroring connection on the source has been disconnected (right-click the connection in the Storage Mirroring Management Console and select **Disconnect**).

13. To begin the restoration process, open the Storage Mirroring Management Console and select **Tools, Restoration Manager**.

 **NOTE:** You can also run the Storage Mirroring DTCL automated restoration script, which can be found in the *HP OpenView Storage Mirroring user's guide*, to complete the remaining steps in this section.

14. Complete the appropriate fields as described below.

- **Original Source**—The name of the source machine where the data originally resided.
- **Restore From**—The name of the target machine that contains the replicated data.
- **Replication Set**—The name of the replication set to be restored.
- **Restore To**—The name of the machine where the data will be restored. This may or may not be the same as the original source machine.

15. Identify the correct drive mappings for the data and any other restoration options necessary. For detailed information on the restoration options, see the *HP OpenView Storage Mirroring user's guide*.

16. Select the **Use block checksum comparison/delta block transfer** option.


17. On the Orphans tab, select **Delete Orphaned Files**.

18. Verify that the selections you have made are correct and click **Restore**. The restoration procedure time will vary depending on the amount of data that you have to restore.

19. After the restoration is complete, start the SQL services on the source machine.

20. Reestablish the Storage Mirroring SQL replication set connection.

21. On the Mirroring tab, select **File differences** with the **Use block checksum** option.

 **NOTE:** If the target has been previously mirrored to or restored, select **File differences** with the **Use block checksum** option so that only the changed data is sent across the network.

At this time, your data is restored back to your source machine, the source machine is again the primary SQL Server, and, if you selected to continue failover monitoring, the target is available to stand in for the source in the event of a failure.

Cluster-to-standalone failover

This section describes an example of how to configure Storage Mirroring and SQL Server. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Requirements

For this scenario, your systems must meet the following requirements:

- Microsoft Windows Server 200x with the latest service pack
- Licensed copies of Microsoft SQL Server 2000 with the latest service pack
- Licensed copies of Storage Mirroring 4.x with the latest service pack


Install Storage Mirroring on the source cluster

1. Install Storage Mirroring on the first node of the cluster. See the *HP OpenView Storage Mirroring getting started guide* for details.
2. Once Storage Mirroring has been successfully installed and the first node is back online, install Storage Mirroring on the second node of the cluster.


Configure the source cluster

The following steps should be completed on the first node of the cluster.

1. If a suitable cluster resource group does not already exist, create a SQL group.
 - a. Right-click the Groups folder on the left pane of the Cluster Administrator and select **New, Group**.
 - b. Specify the Name and Description and click **Next** to continue.

 **NOTE:** The name should be the same as the one you plan on using for the SQL virtual server name. If it is not, the SQL installation will create a cluster group using the SQL virtual server name that you specify during the installation and move all of the SQL resources, including the Physical Disk resource, to that group.

- c. Specify any Preferred Owners and click **Finish** to complete the creation of the new group.

 **NOTE:** You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

2. Create a SQL Physical Disk resource.
 - a. Right-click the SQL group created in step 1 and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name that identifies the SQL shared disk that will contain the databases and log files that you are protecting. This name must be unique within the cluster.
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify **Physical Disk**.
 - **Group**—The SQL Group should be selected. If it is not, select the correct group name.
 - c. Click **Next** to continue.
 - d. The shared Physical Disk resource ensures that the data is available to all nodes listed in the Possible owners list. The default includes all nodes and should not be changed. Click **Next** to continue.
 - e. The shared Physical Disk resource is not dependent on any other resources. Therefore, the default of no resources appearing in the Dependencies dialog box does not need to be changed. Click **Next** to continue.
 - f. Specify the shared Physical Disk parameters using the settings below.

- **Disk**—Select the disk where the SQL data will be stored.
- g. Click **Next** to continue.
- h. Right-click the SQL resource group and select **Bring Online**.
At this point, you should have the SQL resource group established with the shared Physical Disk resources online.

Install SQL Server on the source cluster


Complete these steps on the first node of the cluster.

1. Install SQL Server 2000 following the on-screen instructions.
2. When prompted for the server or instance name, use the group name created in step 1 of "[Configure the source cluster](#)" on page 12.
3. When prompted for the drive to install the virtual SQL Server on, select the shared SQL Physical Disk resource.
4. Apply any SQL 2000 service packs, if applicable.

After you have completed the SQL 2000 installation, you should have the SQL resource group established with SQL installed and all of the resources online.

Install and configure software on the target server

1. Install Storage Mirroring on the target server using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.
2. Install SQL 2000 on the target server using the same drive and directory specifications used on the source cluster. See "[Install SQL Server on the source cluster](#)" on page 13.
3. Select **Start, Programs, Administrative Tools, Services**.
4. Stop the following services and set them to manual startup so that all SQL files are closed on the target and the Storage Mirroring source can replicate the changes. Make sure that you set all of the services, SQL included, relevant to your environment.
 - Distributed Transaction Coordinator
 - Message Queuing
 - MSSQLServer
 - SQLServerAgent
 - Microsoft Search

 **NOTE:** If a failure should occur, the failover and failback scripts that you will be creating will control the stopping and starting of the SQL services on the source and target.

5. Using the Network Connections applet, add a unique IP address to the NIC that will be used for Storage Mirroring mirroring and replication traffic.

Configure mirroring and replication


1. On one of the nodes of the source cluster, select **Start, Programs, Storage Mirroring, Management Console**.
2. Double-click the node that owns the SQL resource group to log on.
3. Right-click the node and select **Properties**.
4. On the Setup tab, disable **Automatically Reconnect on Source Initialization**.
5. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
6. Repeat steps 3, 4, and 5 on the second node of the cluster.
7. Right-click the node that owns the SQL resource group and select **New, Replication Set**. Enter the desired name for the replication set, such as "SQLProtect".

Replication set name: _____

8. Select the SQL data you wish to protect. For more information about what you should include in your SQL replication set, see <cross reference>Configuring the replication set on page 4.
9. Right-click the replication set name and select **Save** to save the replication set.
10. Right-click the replication set name and select **Properties**.
11. Select the Rules tab. Write down the exact drives and paths shown, including the Include/Exclude/Recursive attributes of each.

Path	Inc/Exc/Rec Attributes


12. Right-click the second node of the cluster and select **New, Replication Set**. Enter the *exact* name, including case, that you entered for the first node in step 7.
13. Right-click the replication set name and select **Save** to save the replication set.
14. Right-click the replication set name and select **Properties**.
15. For each drive and path written down in step 11, click **Add** and enter the *exact* path and **Include/Exclude/Recursive** attributes.

 **NOTE:** The replication sets on all nodes must be identical.

Create the Storage Mirroring source connection and begin mirroring and replication


When Storage Mirroring is installed on a cluster, it registers the Storage Mirroring Source Connection resource with the cluster service. The Storage Mirroring Source Connection resource allows Storage Mirroring connections to be made from the cluster to other Storage Mirroring servers.

1. Select **Start, Programs, Administrative Tools, Cluster Administrator**.
2. Right-click the SQL resource group and select **New, Resource**.
3. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name that indicates this is the Storage Mirroring Source Connection.
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify **Storage Mirroring Source Connection**.
 - **Group**—The SQL resource group name should be selected. If it is not, select the correct group name.
4. Click **Next** to continue.
5. Verify that both nodes appear as Possible Owners and click **Next** to continue.
6. This resource is dependent on the disk(s) that contain the data that was included in your replication set. Set the Resource Dependencies to include your disk resource(s) and click **Next** to continue.
7. Specify the following on the Storage Mirroring Source Connection Parameters dialog box:
 - **Replication Set**—Specify the name of the SQL replication set. This name is case-sensitive and should be the same name as specified in "[Configure mirroring and replication](#)" on page 13.
 - **Storage Mirroring Target**—Specify the IP address to be used for Storage Mirroring mirroring and replication (created in step 5 of "[Install and configure software on the target server](#)" on page 13).
 - **Allow connection scripts to interact with the desktop**—Select this check box if you want to display the connection information in a command prompt dialog box.

 **NOTE:** If the connection is set to interact with the desktop, the results will be displayed on the owning node. The resource will remain at the `Online pending` status until the script is manually acknowledged by a user.

Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves. Therefore, this option should not be a permanent setting, but should be used for troubleshooting if you are having problems establishing a connection.

-
- **Logon to Target as**—Allows you to specify an ID, password, and domain to be used when logging on to the target.

 **NOTE:** The user must be a member of the Storage Mirroring Admin security group on all nodes of the cluster running a Storage Mirroring Source Connection resource. If you do not specify a logon, the ID used to start the cluster service will be used. Verify that the ID has Storage Mirroring Admin access.

8. Click **Finish** to complete the creation of the Storage Mirroring Source Connection resource.

9. Right-click the Storage Mirroring Source Connection and choose **Properties**.

10. Click on the Advanced tab. Under Restart, uncheck **Affect the Group**. Click **OK**.

11. Perform this procedure to recover disk space on the target in the event that databases are deleted from the source while the Storage Mirroring Source Connection is offline, leaving orphaned database files on the target. You will need to modify one of the scripts that controls the Storage Mirroring Source Connection resource to enable orphan file removal.

a. Open the file `online.dtc1`, located in the directory where you installed Storage Mirroring, with a text editor.

b. Locate the following section in the file.

```
if ($exitcode = 0) then
    WRITE "Starting block checksum mirror";
    $mirror = MIRROR START $conid DIFFERENT,CHECKSUM;
```

c. Add a space and the keyword `ORPHANS` (before the semicolon) to the last line in this section so that it is identical to the following command:


```
$mirror = MIRROR START $conid DIFFERENT,CHECKSUM ORPHANS;
```

d. Save and close the file.

e. Repeat steps a-d on the second node of the source cluster.

12. Right-click on the Storage Mirroring Source Connection resource and select **Bring Online**.

The Storage Mirroring Source Connection should come online and begin mirroring and replicating the source cluster SQL data to the target server. You can monitor the connection through the Storage Mirroring Management Console.

 **NOTE:** If you start SQL Server and mount the replicated databases on the target, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. This can be accomplished by taking the Storage Mirroring Source Connection resource offline and then bringing it back online.

If the updated data on the target is needed, restore the data from the target to the source.

Set SPN permissions

ServicePrincipalName (SPN) is a property of Active Directory computer and user accounts. The Storage Mirroring service must have the appropriate permissions to modify SPNs on the source cluster's SQL Server

service account. Accordingly, the `WriteServicePrincipalName` permission on the SQL Server's service account in Active Directory must be assigned to the Storage Mirroring service account.

The following procedure includes references to the SQL Server service account and the Storage Mirroring service account. Use the following guidelines to determine which specific Active Directory user or computer account to use:

- If a service is using a user account, then that user account should be used when the service account is referenced.
- If a service is using the Local System account, then the computer account should be used when the service account is referenced.

To find out which account a service is using, run the Windows Services snap-in and view the Log On tab on the service's Properties dialog box to see if the service is using a user account or the Local System account.

Use the following procedure to assign the SQL Server service's `WriteServicePrincipalName` permission to the Storage Mirroring service account. From your domain controller:


1. Click **Start, Run**, and type `MMC`. Click **OK**.
2. Select **File, Add/Remove Snap-in**.
3. Click **Add** and select **ADSI Edit**.
4. Click **Add**. Click **Close**, then click **OK**.
5. Right-click **ADSI Edit** and select **Connect To**.
6. Type in your domain name and click **OK**.
7. Expand the tree in the left pane to the `CN=Users` folder.
8. In the Users folder, locate the source cluster's SQL service account.
9. Right-click on the account and select **Properties**.
10. On the Security tab, click the **Advanced** button.
11. Click **Add**, then click **Object Types**.
12. Select any unchecked check boxes and click **OK**.
13. Type in or locate the target server's Storage Mirroring service account and click **OK**.
14. On the Properties tab, select the **WriteServicePrincipalName** property checkbox.
15. Click **OK** to close all open dialog boxes, then click **Close** to close the window.

Configure failover and begin failure monitoring

Perform the following steps on the target server.

1. If the SQL Server service is **not** using the same domain user account on the source and target servers, use the NSISPN utility to modify the SQL SPNs of the source and target servers when failover and failback occur. Create a batch file called `prefailover.bat` using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

NOTE: If you **are** using the same domain user account on the source and target servers, you do not need to update SPN permissions. Proceed to the next step

-  After you modify the sample scripts, save them with a new name to remove the `sample_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

sample_prefailover.bat

```
rem ***SAMPLE*** pre-failover script for cluster-to-standalone failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem When the SQL Server service is configured to use different domain user accounts
rem on the source and target servers, use commands similar to the following.
rem The default SQL Server port, 1433, is used in these examples.
rem SourceSqlHost refers to the SQL virtual server name on the source cluster.

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlServiceAccount
NSISPN -D MSSQLSvc/SourceSqlHost:1433 SourceSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost:1433 TargetSqlServiceAccount

rem When the SQL Server service is started with the server's System account
rem use commands similar to the following.

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlHost
NSISPN -D MSSQLSvc/SourceSqlHost:1433 SourceSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost:1433 TargetSqlHost
```

2. If a failure occurs, you will want the SQL services to start on the target machine automatically. To do this, create a batch file called `postover.bat` using the sample batch file below. The following script should be run after failover has completed.

sample_postover.bat

```
rem ***SAMPLE*** post-failover script for cluster-to-standalone failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem You may need to remark out some commands depending on the function of your target.
rem The following commands should be used if your target is functioning as a back-end
rem Make sure that you start all of the services, SQL included, relevant to your
rem environment.

net start "Distributed Transaction Coordinator"
net start "MSSQLSERVER"
net start "SQLSERVERAGENT"
net start "Microsoft Search"
```

3. After a cluster failure is resolved, you will be ready to bring your source back online. At that time, you will want to stop the SQL services on the target automatically. To do this, create a batch file called

preback.bat using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

If you are using the same domain user account on the source and target servers, you do not need to update SPN permissions. Remove the NSISPN commands in the sample script.

sample_preback.bat

```
rem ***SAMPLE*** pre-failback script for cluster-to-standalone failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem You may need to remark out some commands depending on the function of your target.
rem The following commands should be used if your target is functioning as a back-end.
rem Make sure that you start all of the services, SQL included, relevant to your
rem environment.


net stop "Distributed Transaction Coordinator"
net stop "MSSQLSERVER" /y
net stop "SQLSERVERAGENT"
net stop "Microsoft Search"

rem When the SQL Server service is configured to use different domain user accounts
rem on the source and target servers, use commands similar to the following.
rem The default SQL Server port, 1433, is used in these examples
rem SourceSqlHost refers to the SQL virtual server name on the source cluster.


NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlServiceAccount
NSISPN -D MSSQLSvc/SourceSqlHost:1433 TargetSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlServiceAccount
NSISPN -A MSSQLSvc/SourceSqlHost:1433 SourceSqlServiceAccount

rem When the SQL Server service is started with the server's System account
rem use commands similar to the following in the failback scripts:

NSISPN -D MSSQLSvc/SourceSqlHost.domain.com:1433 TargetSqlHost
NSISPN -D MSSQLSvc/SourceSqlHost:1433 TargetSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost.domain.com:1433 SourceSqlHost
NSISPN -A MSSQLSvc/SourceSqlHost:1433 SourceSqlHost
```

 **NOTE:** The Windows Resource Kit SETSPN utility may be used in place of the NSISPN utility.

4. On the target, select **Start, Programs, Storage Mirroring, Failover Control Center**.
5. Select the target machine from the list of available machines. If the target you need is not displayed, click **Add Target**, enter the machine name, click **OK**, and then click **Login**.
6. To add a monitor for the selected target, click **Add Monitor**. Type the name of the SQL source virtual server and click **OK**. The Monitor Settings window will open.
7. In the left pane of the Monitor Settings window, mark *only* the SQL source virtual server IP address that is going to failover. Verify that the correct NIC is selected in the right-hand pane under Current IP Address(es)."
8. Configure the Monitor Interval, Missed Packets, and Timeout.
 - **Monitor Interval**—The monitor interval specifies how often, in seconds, the monitor request is sent to the source machine. The default is 5 seconds.
 - **Missed Packets**—Missed packets specifies how many monitor replies can be missed before assuming the source machine has failed. The default is 5 replies.
 - **Timeout**—The failover timeout is the amount of time before failover begins. The time is calculated by multiplying the values of the monitor interval and missed packets settings. This time is displayed on the Failover Control Center.

 **NOTE:** Make sure that the source timeout is long enough to allow for a node-to-node movement of the SQL resource group on the source cluster. The default timeout of 25 seconds (5 seconds X 5 missed packets) may not allow enough time for the SQL resource group to move from node to node without a failover occurring. Move the SQL group and time how long the SQL virtual IP address resource is offline before coming online on the other node. Adjust the Monitor Interval and Missed Packets settings as needed. A Monitor Interval of 6 seconds and a Missed Packets setting of 10 will yield 60 seconds of Timeout before failover occurs.

9. Highlight the SQL source virtual server name under Names to Monitor and click **Scripts**. Specify the scripts that were created earlier using `prefailover.bat` for the target pre-failover script, `postover.bat` for the target post-failover script, and `preback.bat` for the target pre-failback script.
10. Under Items to Failover, verify that only Monitored IP Addresses and Server Name are selected.
11. Click **OK** to go back to the Monitor Settings dialog box and then click **OK** to begin monitoring the source machine.

In the event of a source machine failure, your target machine is now ready to stand in for your source. For information on monitoring failover, see the *HP OpenView Storage Mirroring user's guide*.

Dealing with a failure

In the event that the source cluster or source SQL virtual server becomes unavailable, the following will occur:

1. The Failover Control Center will count down to zero, and a dialog box will appear on the target server indicating the failure. You may initiate failover by clicking **OK**.
2. After clicking OK, the target server will assume the failover IP address of the source cluster SQL virtual server and the server name.
3. The failover scripts will execute, making the necessary changes to the SPNs and starting the SQL services on the target server.
4. Use `NSISPN -L SourceSqlHost` and `NSISPN -L TargetSqlHost` to verify that the changes were made.

 **NOTE:** The Windows Resource Kit SETSPN utility can also be used in place of the NSISPN utility.

The target server is now standing in for the source SQL virtual server. Clients may access SQL services on the target just as if it were the original source.

Fail back to the source cluster

To fail back to your source cluster, you will need to complete the following:

- "Update the source with the new data from the target" on page 19
- "Fail back to the source cluster" on page 20

Update the source with the new data from the target

During the source downtime, users are updating data on the target server. When your source cluster is ready to come back online, the data on the source cluster is no longer current and must be updated with the new data from the target server.

1. On the target server, use the Network Connection applet to remove the IP address used for Storage Mirroring mirroring and replication. This will prevent the source cluster from overwriting the target data when the source comes online.

 **NOTE:** If a separate NIC is being used for Storage Mirroring traffic, you may disable it rather than removing the IP address.

2. Bring one or both of the source cluster nodes online.

3. Using Cluster Administrator, bring all source cluster SQL resources offline.
4. On the target server, select **Start, Programs, Storage Mirroring, Management Console**.
5. Double-click the target server to log on.
6. Right-click the target server and select **New, Replication Set**. Enter a name for the replication set, such as "SQLRestore".
7. Select the SQL data you wish to protect. This replication set should be identical to the replication set(s) created on the source cluster. For more information about what you should include in your SQL replication set, see <cross reference>Configuring the replication set on page 4.
8. Right-click the replication set name and select **Save** to save the replication set.
9. On the source cluster, bring only the SQL Physical Disk resources online. Take note of which cluster node owns the SQL resource group.
10. From the Storage Mirroring Management Console on the target server, click on the replication set created in step 6 and drag it onto the source cluster node that owns the SQL resource group.
11. In the Connection Manager, verify that the Source Server specified is your target server, the Replication Set is the one created on the target server, and the Target Server specified is the node that owns the SQL resource group. In the Mappings section, click **One to One**.
12. On the Orphans tab, select the **Move/Delete Orphan Files** box, then click **Connect** to start mirroring and replication from the target server back to the source cluster.
13. Monitor the mirroring process using the Storage Mirroring Management Console to ensure that it completes with no errors. The mirroring process is complete when the Mirror Status is Idle.

Clients may continue to access SQL on the target server while updating the source cluster. Any changes that occur during this process will be replicated from the target server to the source cluster.

After the mirror completes, schedule downtime to convert operations back to the source cluster.


Fail back to the source cluster

At the scheduled downtime, you will be taking the SQL resources offline on the target, removing the SQL source virtual server IP address and name from the target, and deleting a registry key.

1. From the Failover Control Center on the target server, click the **Failback** button. The `preback.bat` file will execute and stop the SQL services on the target.

 **NOTE:** Clients will be unable to access SQL at this time.

2. Failback is complete when the Continue/Stop Monitoring window appears. **Do not choose Continue at this time.**
3. Verify that all data possibly queued on the target server has been sent to the source cluster before continuing. You can verify this using the Bytes in Disk Queue, Bytes in Replication Queue, and Bytes in Mirror Queue statistics from the Storage Mirroring Connection object of Performance Monitor.
4. Verify that all of the data in queue on the source cluster has been written to disk before continuing. You can verify that the queue is empty by checking the Bytes in Target Disk Queue statistic in the Target section of DTStat or the Bytes in Queue statistic in the Storage Mirroring Target section of Performance Monitor. If these statistics are zero (0), the queue is empty and you can continue. If these statistics are not zero, there is still data in queue on the source cluster and you must wait before continuing. (For information on DTStat and Performance Monitor statistics, see the *HP OpenView Storage Mirroring user's guide*.)
5. In the Storage Mirroring Management Console, double-click the target server to log in. In the right pane, right-click on the connection to the source cluster and select **Disconnect**.
6. Start REGEDIT on the target server and browse to `HKEY_LOCAL_MACHINE\Software\NSI Software\Failover`. Delete the `SourceReplicaState` key.

 **NOTE:** Any other source servers failed over to this target will have their Restore Required flags cleared by deleting the `SourceReplicaState` key.

7. As done in "Install and configure software on the target server" on page 13, use the Network Connections applet to add the IP address that was removed in section "Update the source with the new data from the target" on page 19 back to the NIC that will be used for Storage Mirroring mirroring and replication traffic.

 **NOTE:** If you use a separate NIC for Storage Mirroring traffic and disabled it rather than removed the IP address, re-enable the NIC at this time.

8. Using Cluster Administrator, bring the source cluster SQL resource group online. Clients may now connect and use the SQL services on the source cluster. Any changes to source SQL data will be replicated to the target.
9. After the SQL virtual server on the source cluster is online, click **Continue** in the Failover Control Center to continue monitoring the source SQL virtual server for failure.

Cluster-to-cluster failover

This section describes an example of how to configure Storage Mirroring and SQL Server. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Requirements

For this scenario, your system must meet the following requirements:

- Microsoft Windows Server 200x with the latest service pack
- Licensed copies of Microsoft SQL Server 2000 with the latest service pack
- Licensed copies of Storage Mirroring 4.x with the latest service pack

Install Storage Mirroring on the source cluster


Install Storage Mirroring on the first node of the source cluster. See the *HP OpenView Storage Mirroring getting started guide* for details.

Once Storage Mirroring has been successfully installed and the first node is back online, install Storage Mirroring on the second node of the source cluster.


Configure the source cluster

The following steps should be completed on the first node of the source cluster.

1. If a suitable cluster resource group does not already exist, create a SQL group.
 - a. Right-click the Groups folder on the left pane of the Cluster Administrator and select **New, Group**.
 - b. Specify the Name and Description and click **Next** to continue.

 **NOTE:** The name should be the same as the one you plan on using for the SQL virtual server name. If it is not, the SQL installation will create a cluster group using the SQL virtual server name that you specify during the installation and move all of the SQL resources, including the Physical Disk resource, to that group.

- c. Specify any Preferred Owners and click **Finish** to complete the creation of the new group.

 **NOTE:** You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

2. Create a SQL Physical Disk resource.
 - a. Right-click the SQL group created in step 1 and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:

- **Name**—Specify a name that identifies the shared disk that will contain the databases and log files that you are protecting. This name must be unique within the cluster.
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify Physical Disk.
 - **Group**—The SQL Group should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
 - d. The shared Physical Disk resource ensures that the data is available to all nodes listed in the Possible owners list. The default includes all nodes and should not be changed. Click **Next** to continue.
 - e. The shared Physical Disk resource is not dependent on any other resources. Therefore, the default of no resources appearing in the Dependencies dialog box does not need to be changed. Click **Next** to continue.
 - f. Specify the shared Physical Disk parameters using the settings below.
 - **Disk**—Select the disk where the SQL data will be stored.
 - g. Click **Next** to continue.
 - h. Right-click the SQL resource group and select **Bring Online**.

At this point, you should have the SQL resource group established on the source cluster with the shared Physical Disk resource online.

Install SQL Server on the source cluster

Complete these steps on the first node of the source cluster.


1. Install SQL Server 2000 following the on-screen instructions.
2. When prompted for the server or instance name, use the group name created in step 1 of "[Configure the source cluster](#)" on page 21.
3. When prompted for the drive to install the virtual SQL Server on, select the shared SQL Physical Disk resource.
4. Install any SQL 2000 service packs, if applicable.

After you have completed the SQL 2000 installation on the source cluster, you should have the SQL resource group established with SQL installed and all of the resources online.

Install Storage Mirroring on the target cluster


1. Install Storage Mirroring on the first node of the target cluster. See the *HP OpenView Storage Mirroring getting started guide* for details.
2. Once Storage Mirroring has been successfully installed and the first node is back online, install Storage Mirroring on the second node of the target cluster.

Configure the target cluster


 **NOTE:** Although you will be configuring a unique virtual server on the target cluster, the target cluster physical drives used for SQL databases and log files must have identical drive letters to the source cluster.

The following steps should be completed on the first node of the target cluster.

1. If a suitable cluster resource group does not already exist, create a SQL group.
 - a. Right-click the Groups folder on the left pane of the Cluster Administrator and select **New, Group**.
 - b. Specify the Name and Description and click **Next** to continue.

 **NOTE:** The name should be the same as the one you plan on using for the SQL virtual server name, but it must be unique and not the same name as was used on the source cluster. If you enter a name that is different from the name you plan to use for the SQL virtual server, the SQL installation will create a cluster group using the SQL virtual server name that you specify during the installation and move all of the SQL resources, including the Physical Disk resource, to that group.

- c. Specify any Preferred Owners and click **Finish** to complete the creation of the new group.


 **NOTE:** You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

2. Create a SQL Physical Disk resource.

- a. Right-click the SQL group created in step 1 and select **New, Resource**.
- b. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name that identifies the shared disk that will contain the databases and log files that you are protecting. This name must be unique within the cluster.
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify Physical Disk.
 - **Group**—The SQL Group should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
- d. The shared Physical Disk resource ensures that the data is available to all nodes listed in the Possible owners list. The default includes all nodes and should not be changed. Click **Next** to continue.
- e. The shared Physical Disk resource is not dependent on any other resources. Therefore, the default of no resources appearing in the Dependencies dialog box does not need to be changed. Click **Next** to continue.
- f. Specify the shared Physical Disk parameters using the settings below.
 - **Disk**—Select the disk where the SQL data will be stored.
- g. Click **Next** to continue.
- h. Right-click the SQL resource group and select **Bring Online**.

At this point, you should have the SQL resource group established on the target cluster with the shared Physical Disk resource online.

Install SQL Server on the target cluster

 **NOTE:** Although you will be configuring a separate SQL virtual server on the target cluster, the target cluster physical drives and paths used for SQL databases and log files must be identical to the source cluster.

Complete these steps on the first node of the target cluster.

1. Install SQL 2000 following the on-screen instructions.
2. When prompted for the server or instance name, use the group name created in step 1 of "[Configure the target cluster](#)" on page 22.
3. When prompted for the drive to install the virtual SQL Server on, select the shared Physical Disk resource.
4. Install any SQL 2000 service packs, if applicable.
5. Create an IP address for Storage Mirroring replication and mirroring.
 - a. Using Cluster Administrator, right-click on the target SQL group and choose **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name, such as "Storage Mirroring IP Address"
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify IP Address.
 - **Group**—Select the target SQL group.
 - c. Click **Next**.
 - d. Verify that both target nodes appear as Possible Owners and click **Next** to continue.
 - e. IP Address resources are not dependent on any other resources. Click **Next** to continue.
 - f. Provide an IP address and a network mask, and select the network to be used for Storage Mirroring traffic.
 - g. Click **Finish**, then click **OK**.
 - h. Right-click on the newly-created IP address and select **Bring Online**.

After you have completed the SQL 2000 installation on the target cluster, you should have the SQL resource group established with SQL installed and all of the resources online.


Configure mirroring and replication

1. On one of the nodes of the source cluster, select **Start, Programs, Storage Mirroring, Management Console**.
2. In the left pane, double-click the node that owns the SQL resource group to log on.
3. Right-click on the node and select **Properties**.
4. On the Setup tab, disable **Automatically Reconnect on Source Initialization**.
5. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
6. Double-click on the second node of the source cluster to log in.
7. Repeat steps 3, 4, and 5 on the second node of the source cluster.
8. Right-click the node that owns the SQL resource group and select **New, Replication Set** and enter the desired name for the replication set, such as "SQLProtect".
9. Select the SQL data you wish to protect. For more information about what you should include in your SQL replication set, see "[Configuring the replication set](#)" on page 4 .
10. Right-click the replication set name and select **Save** to save the replication set.
11. Right-click the replication set name and select **Properties**.

12. Select the Rules tab. Write down the exact drives and paths shown, including the Include/Exclude/Recursive attributes of each.

Path	Inc/Exc/Rec Attributes

13. Right-click the second node of the cluster and select **New, Replication Set** and enter the exact name, including case, that you entered for the first node.
14. Right-click the replication set name and select **Save** to save the replication set.
15. Right-click the replication set name and select **Properties**.
16. For each drive and path written down in step 12, click **Add** and enter the exact path and Inc/Exc/Rec attributes.


 **NOTE:** The replication sets on each node of the source cluster must be identical to each other.

Create the Storage Mirroring Source Connection and begin mirroring and replication

When Storage Mirroring is installed on a cluster, it registers the Storage Mirroring Source Connection resource with the cluster service. The Storage Mirroring Source Connection resource allows Storage Mirroring connections to be made from the cluster to other Storage Mirroring servers.


1. On a *target cluster node*, select **Start, Programs, Administrative Tools, Cluster Administrator**.
2. Right-click the target cluster SQL resource group and select **Take Offline**.
3. Within the target SQL resource group, right-click the SQL virtual IP address and select **Bring Online**. This provides a target cluster SQL IP address to mirror and replicate to from the source cluster.
4. Right-click the Physical Disk resource and select **Bring Online**. This provides a target cluster SQL Physical Disk to mirror and replicate to from the source cluster.
5. On a *source cluster node*, select **Start, Programs, Administrative Tools, Cluster Administrator**.
6. On the *source cluster*, right-click on the SQL resource group and select **New, Resource**.
7. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name that indicates this is the Storage Mirroring source connection.
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify Storage Mirroring Source Connection.
 - **Group**—The SQL resource group name should be selected. If it is not, select the correct group name.
8. Click **Next** to continue.
9. Verify that both nodes appear as Possible Owners and click **Next** to continue.
10. This resource is dependent on the disk that contains the data that was included in your replication set. Set the Resource Dependencies to include your disk resource(s) and click **Next** to continue.
11. Specify the following on the Storage Mirroring Source Connection Parameters dialog box:
 - **Replication Set**—Specify the name of the SQL replication set created on the source cluster nodes. This name is case-sensitive and must be the same name as specified in Configure mirroring and replication on page 31.

- **Storage Mirroring Target**—Specify the Storage Mirroring IP address created within the target SQL group in step 5 of “[Install SQL Server on the target cluster](#)” on page 24.
- **Allow connection scripts to interact with the desktop**—Enable this check box if you want to display the connection information in a command prompt dialog box.

 **NOTE:** If the script is set to interact with the desktop, the results will be displayed on the owning node. The resource will remain at the Online pending status until the script is manually acknowledged by a user.

Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves. Therefore, this option should not be a permanent setting, but should be used for troubleshooting if you are having problems establishing a connection.

-
- **Logon to Target as**—Allows you to specify an ID, password, and domain to be used when logging on to the target.

 **NOTE:** The user must be a member of the Storage Mirroring Admin security group on all nodes of the cluster running a Storage Mirroring Source Connection resource. If you do not specify a logon, the ID used to start the cluster service will be used. Verify that the ID has Storage Mirroring Admin access.

12. Click **Finish** to complete the creation of the Storage Mirroring Source Connection resource.

13. Right-click the Storage Mirroring Source Connection and choose **Properties**.

14. Click on the Advanced tab. Under Restart, uncheck **Affect the Group**. Click **OK**.

15. Perform this procedure to recover disk space on the target in the event that databases are deleted from the source while the Storage Mirroring Source Connection is offline, leaving orphaned database files on the target. You will need to modify one of the scripts that controls the Storage Mirroring Source Connection resource to enable orphan file removal.

a. Open the file `online.dctl`, located in the directory where you installed Storage Mirroring, with a text editor.

b. Locate the following section in the file.

```
if ($exitcode = 0) then
    WRITE "Starting block checksum mirror";
    $mirror = MIRROR START $conid DIFFERENT,CHECKSUM;
```

c. Add a space and the keyword `ORPHANS` (before the semicolon) to the last line in this section so that it is identical the following command:


```
$mirror = MIRROR START $conid DIFFERENT,CHECKSUM ORPHANS;
```

d. Save and close the file.

e. Repeat steps a-d on the second node of the source cluster.

16. Right-click on the Storage Mirroring Source Connection resource and select **Bring Online**.

The Storage Mirroring Source Connection should come online and begin mirroring and replication of the source cluster SQL data to the target server. You can monitor the connection through the Storage Mirroring Management Console.

 **NOTE:** If you start SQL Server and mount the replicated databases on the target cluster, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. This can be accomplished by taking the Storage Mirroring Source Connection resource offline, and then bringing it back online.

If the updated data on the target is needed, restore the data from the target to the source.

Set SPN permissions

ServicePrincipalName (SPN) is a property of Active Directory computer and user accounts. The Storage Mirroring service must have the appropriate permissions to modify SPNs on the source cluster's SQL Server service account. Accordingly, the `WriteServicePrincipalName` permission on the SQL Server's service account in Active Directory must be assigned to the Storage Mirroring service account.

The following procedure includes references to the SQL Server service account and the Storage Mirroring service account. Use the following guidelines to determine which specific Active Directory user or computer account to use:

- If a service is using a user account, then that user account should be used when the service account is referenced.
- If a service is using the Local System account, then the computer account of each node should be used when the service account is referenced.

To find what account a service is using, run the Windows Services snap-in and check the Log On tab on the service's Properties dialog to see if the service is using a user account or the Local System account.

Use the following procedure to assign the SQL Server service's `WriteServicePrincipalName` permission to the Storage Mirroring service account. From your domain controller:

1. Click **Start, Run**, and type MMC. Click **OK**.
2. Select **File, Add/Remove Snap-in**.
3. Click **Add** and select **ADSI Edit**.
4. Click **Add**. Click **Close**, then click **OK**.
5. Right-click **ADSI Edit** and select **Connect To**.
6. Type in your domain name and click **OK**.
7. Expand the tree in the left pane to the CN=Users folder.
8. In the Users folder, locate the source cluster's SQL service account.
9. Right-click on the account and select **Properties**.
10. On the Security tab, click the **Advanced** button.
11. Click **Add**, then click **Object Types**.
12. Select any unchecked boxes and click **OK**.
13. Type in or locate a target node's Storage Mirroring service account and click **OK**.
14. On the Properties tab, select the `WriteServicePrincipalName` property.
15. Repeat steps 13 and 14 for the other target node.
16. Click **OK** to close all open dialog boxes, then click **Close** to close the window.

Dealing with a failure

In the event the source cluster should fail, you will want the target cluster to bring its SQL resource group online. Prior to this, the source's SQL virtual IP address resource should be added to the target SQL resource group.

The following steps should be performed on the target machine if a failure of the source cluster occurs:


1. On a target node, select **Start, Programs, Administrative Tools, Cluster Administrator**.
2. Right-click the SQL resource group and select **New, Resource**.
3. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name for the resource, such as "Duplicate SQL Source IP".
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify IP Address.
 - **Group**—The SQL resource group name should be selected. If it is not, select the correct group name.
4. Click **Next** to continue.
5. Verify that both target cluster nodes appear as Possible Owners and click **Next** to continue.

6. The IP Address resource is not dependant upon any other resources, so click **Next** to continue.
7. Provide the source SQL virtual IP address and the subnet mask, and select the network that clients will use to connect to the SQL virtual server.
8. Click **Finish**, then click **OK**.
9. Right-click the SQL Network Name resource and select **Properties**.
10. Select the Dependencies tab and click **Modify**.
11. Add the duplicate source IP address to the Dependencies list and click **OK**, and **OK** again.
12. Right-click the SQL resource group and select **Bring Online**.

The SQL resource group will come online and clients may now access the SQL virtual server.

The process of creating the duplicate source IP on the target cluster can be automated using the cluster command line interface. The batch file can be kept on a target node until needed. For more information on using the cluster command line interface, see your Windows documentation or type "Cluster /?" at a command prompt.

The following sample batch file can be used to automate steps 1-12.

 **NOTE:** After you modify the sample scripts, save them with a new name to remove the `sample_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment

sample_failover.bat

```
rem ***SAMPLE*** failover script for cluster-to-cluster failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem "TargetCluster"- The actual name of the target cluster.
rem "SQLServerBak" - The actual name of the TARGET cluster SQL resource group.
rem "Client LAN" - The name of the network that clients uses to connect to the TARGET cluster.
rem Address=169.254.9.120 - The SOURCE SQL virtual server IP address.
rem SubnetMask=255.255.0.0 - The SOURCE SQL virtual server subnet mask.
rem Description - Optional, but it is less confusing if it is the SOURCE SQL virtual server IP address.
rem "SQL Network Name(SQLServer)" - The actual NAME of the network name resource as it appears in
rem Cluster Administrator.

rem Create a duplicate source IP within the target cluster SQL group
cluster /cluster:"TargetCluster" res "Duplicate SQL Source IP" /create /group:"SQLServerBak" /type:"IP Address"
cluster /cluster:"TargetCluster" res "Duplicate SQL Source IP" /priv Network="Client LAN"
cluster /cluster:"TargetCluster" res "Duplicate SQL Source IP" /priv Address=169.254.9.120
cluster /cluster:"TargetCluster" res "Duplicate SQL Source IP" /priv SubnetMask=255.255.0.0
cluster /cluster:"TargetCluster" res "Duplicate SQL Source IP" /prop Description="169.254.9.120"
cluster /cluster:"TargetCluster" res "SQL Network Name(SQLServer)" /adddep:"Duplicate SQL Source IP"

rem Bring target cluster SQL group online
cluster /cluster:"TargetCluster" group "SQLServerBak" /online
```

Fail back to the source cluster

To fail back to your source cluster, you will need to complete the following:

- "Update the source with the new data from the target"
- "Update the source with the new data from the target" on page 29
- "Bring the source SQL group online"

Update the source with the new data from the target

During the source downtime, users are updating data on the target cluster. When your source cluster is ready to come back online, the data is no longer current and must be updated with the new data from the target cluster.

1. On the target cluster, use Cluster Administrator to take the Storage Mirroring IP Address resource offline (created in step 5 of "[Install SQL Server on the target cluster](#)"). This will prevent the Storage Mirroring Source Connection resource on the source cluster from overwriting data on the target cluster when the source comes online.
2. Bring one or both of the source cluster nodes online.
3. Using Cluster Administrator, bring all source cluster SQL resources offline.
4. Using the Storage Mirroring Management Console, double-click the target node in the left pane that owns the SQL resource group to log in.
5. Right-click the node and select **New, Replication Set**. Enter the desired name for the replication set, such as "SQLRestore".
6. Select the SQL data you wish to protect. The replication set should be identical to the replication set(s) created on the source cluster. For more information about what you should include in your SQL replication set, see "[Configuring the replication set](#)" on page 4.
7. Right-click the replication set name and select **Save to save the replication set**.
8. Right-click the replication set name and select **Properties**.
9. Select the Rules tab. Write down the exact drives and paths shown, including the Include/Exclude/Recursive attributes of each.


Path	Inc/Exc/Rec Attributes

10. Double-click on the second node of the target cluster to log in.
11. Right-click the second node of the cluster and select **New, Replication Set** and enter the exact name, including case, that you entered for the first node.
12. Right-click the replication set name and select **Save to save the replication set**.
13. Right-click the replication set name and select **Properties**.
14. For each drive and path written down in step 9, click **Add** and enter the exact path and Inc/Exc/Rec attributes.

 **NOTE:** The replication sets on all nodes must be identical.


15. On a target cluster node, select Start, Programs, Administrative Tools, Cluster Administrator.
16. On the *target cluster*, right-click on the SQL resource group and select **New, Resource**.
17. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify a name that indicates this is the Storage Mirroring source connection.
 - **Description**—You can optionally add a more detailed description for this resource.
 - **Resource type**—Specify Storage Mirroring Source Connection.
 - **Group**—The target SQL resource group name should be selected. If it is not, select the correct group name.
18. Click **Next** to continue.

19. Verify that both nodes appear as Possible Owners and click **Next** to continue.
20. This resource is dependent on the disk that contains the data that was included in your replication set. Verify that the Resource Dependencies include your disk resource and click **Next** to continue.
21. Specify the following on the Storage Mirroring Source Connection Parameters dialog box:
 - **Replication Set**—Specify the name of the SQL replication set created on the target cluster nodes. This name is case-sensitive.
 - **Storage Mirroring Target**—Specify the IP address or name of the node that owns the SQL group on the source cluster. Do not use the source SQL virtual IP address.
 - **Allow connection scripts to interact with the desktop**—Enable this check box if you want to display the connection information in a command prompt dialog box.

 **NOTE:** If the script is set to interact with the desktop, the results will be displayed on the owning node. The resource will remain at the Online pending status until the script is manually acknowledged by a user.

Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves. Therefore, this option should not be a permanent setting, but should be used for troubleshooting if you are having problems establishing a connection.

-
- **Logon to Target as**—Allows you to specify an ID, password, and domain to be used when logging on to the target.

 **NOTE:** The user must be a member of the Storage Mirroring Admin security group on all nodes of the cluster running a Storage Mirroring Source Connection resource. If you do not specify a logon, the ID used to start the cluster service will be used. Verify that the ID has Storage Mirroring Admin access.

-
22. Click **Finish** to complete the creation of the Storage Mirroring Source Connection resource.
 23. Perform this procedure to recover disk space on the target in the event that databases are deleted from the source while the Storage Mirroring Source Connection is offline, leaving orphaned database files on the target. You will need to modify one of the scripts that controls the Storage Mirroring Source Connection resource to enable orphan file removal.
 - a. Open the file `online.dtcl`, located in the directory where you installed Storage Mirroring, with a text editor.
 - b. Locate the following section in the file.

```
if ($exitcode = 0) then
WRITE "Starting block checksum mirror";
$mirror = MIRROR START $conid DIFFERENT,CHECKSUM;
```
 - c. Add a space and the keyword **ORPHANS** (before the semicolon) to the last line in this section so that it is identical to the following command:

```
$mirror = MIRROR START $conid DIFFERENT,CHECKSUM ORPHANS;
```
 - d. Save and close the file.
 - e. Repeat steps a-d on the second node of the target cluster.

24. On the source cluster, right-click the Physical Disk resource within the SQL resource group and select **Bring Online**.


 **NOTE:** Bring only the Physical Disk resource online at this time.

-
25. On the target cluster, right-click the Storage Mirroring Source Connection resource within the SQL resource group and select **Bring Online**.

The Storage Mirroring Source Connection should come online and begin mirroring and replication the target cluster SQL data to the source cluster. You can monitor the connection through the Storage Mirroring

Management Console. Clients may continue to access SQL services on the target cluster. Any changes to the SQL data will be replicated to the source cluster.

Schedule downtime to return operations back to the source cluster.

 **NOTE:** If before the mirror completes, you start the source SQL Server and mount the replicated databases on the source cluster, or if the data on the source is otherwise modified, the data on the source and target will no longer match. If the updated data on the source is not needed, perform a full or difference with block checksum mirror from the target to the source. This can be accomplished by taking the Storage Mirroring Source Connection resource offline and then bringing it back online.

If the updated data on the source is needed, restore the data from the source to the target.


Take the resources offline

At the scheduled downtime, you will be taking the target SQL resources offline, removing a dependency, and deleting two resources.

1. Using Cluster Administrator, take only the SQL resources offline on the target cluster. This can be accomplished by taking the target SQL server virtual IP address offline.

 **NOTE:** Clients will be unable to access SQL at this time.

2. Verify that all data possibly queued on the target cluster has been sent to the source cluster before continuing. You can verify this using the Bytes in Disk Queue, Bytes in Replication Queue, and Bytes in Mirror Queue statistics from the Storage Mirroring Connection object of Performance Monitor.
3. Verify that all of the data in queue on the source cluster has been applied to the source node before continuing. You can verify that the target queue is empty by checking the Bytes in Target Disk Queue statistic in the Target section of DTStat or the Bytes in Queue statistic in the Storage Mirroring Target section of Performance Monitor. If these statistics are zero (0), the queue is empty and you can continue. If these statistics are not zero, there is still data in queue on the source node and you must wait before continuing. (For information on DTStat and Performance Monitor statistics, see the *HP OpenView Storage Mirroring user's guide*.)
4. Once all of the data has been replicated, take the Storage Mirroring Source Connection resource offline on the target cluster and then delete it.
5. Take any other online resources in the target SQL cluster group offline.
6. Remove the duplicate source IP from the target SQL Network Name dependency list, then delete the duplicate source IP.

 **NOTE:** The Storage Mirroring Source Connection resource and the duplicate source IP address on the target cluster are no longer needed after data has been restored to the original cluster.

Bring the source SQL group online

Complete the following steps to bring the source SQL group online.

1. Using Cluster Administrator, bring the target SQL virtual IP address, the Storage Mirroring IP Address created in step 5 of "[Install SQL Server on the target cluster](#)", and the Physical Disk resources online so that the Storage Mirroring Source Connection on the source cluster can connect.

 **NOTE:** Bring **ONLY** the target SQL virtual server IP Address, Storage Mirroring IP Address, and Physical Disk resources online.

2. Bring the source cluster SQL group online.

The failback process can be automated using cluster command line interface commands within a batch file. The batch file can be kept on a target or source node until needed. The following sample batch file can be used to automate failback after the restore mirror has completed.

NOTE:

- The sample batch files provided are only examples and **MUST** be edited to reflect your environment.
 - For more information on using the cluster command line interface, see your Windows documentation or type "Cluster /?" at a command prompt.
-

sample_failback.bat

```
rem ***SAMPLE*** failback script for cluster-to-cluster failover with SQL 2000 servers

rem This sample batch file is provided as an example only. Because no two
rem environments or configurations are exactly the same, you MUST modify
rem this script in order to make the solution work in your environment.

rem "TargetCluster"-The actual name of the target cluster.
rem "SourceCluster"-The actual name of the source cluster.
rem "SQLServerGroup"-The actual name of the source or target cluster SQL resource group.
rem "SQL Network Name(SQLSERVER)"-The actual NAME of the network name resource as it appears in Cluster
rem Administrator.
rem "SQL Data Disk"-The name of the physical disk where the SQL databases are located, as it appears in
rem Cluster Administrator.
rem "SQL IP Address1(SQLSERVER)"-The name of the SQL IP address resource, as it appears in Cluster
rem Administrator.
rem "SQL Restore DTSC"-The name of the Storage Mirroring Source Connection on the target cluster, as it
rem appears
rem in Cluster Administrator.
rem "Storage Mirroring IP Address"-The name of the IP address on the target cluster used for Storage
rem Mirroring
rem mirroring and replication, as it appears in Cluster Administrator.

rem Take the target cluster SQL group offline
cluster /cluster:"TargetCluster" group "SQLServerGroup" /offline

rem Remove duplicate SQL source IP and the DTSC
cluster /cluster:"TargetCluster" res "SQL Network Name(SQLSERVER)" /removedep:"Duplicate SQL Source IP"
cluster /cluster:"TargetCluster" res "Duplicate SQL Source IP" /delete
cluster /cluster:"TargetCluster" res "SQL Restore DTSC" /removedep:"SQL Data Disk"
cluster /cluster:"TargetCluster" res "SQL Restore DTSC" /delete

rem Bring target cluster SQL IP address, Storage Mirroring IP Address, and physical disk resources online
cluster /cluster:"TargetCluster" res "SQL IP Address1(SQLSERVER)" /online
cluster /cluster:"TargetCluster" res "SQL Data Disk" /online
cluster /cluster:"TargetCluster" res "Storage Mirroring IP Address" /online

rem Bring the source cluster MSSQL group online
cluster /cluster:"SourceCluster" group "SQLServerGroup" /online
```

The SQL services will start on the source cluster and the Storage Mirroring Source Connection on the source will begin mirroring back to the target. Clients may now access SQL services on the source cluster. Any changes to the SQL data will be replicated to the target cluster.